First Hit Previous Doc Next Doc Go to Doc#

Generate Collection Print

L4: Entry 1 of 4 File: PGPB Mar 24, 2005

DOCUMENT-IDENTIFIER: US 20050066100 A1
TITLE: System having storage subsystems and a link coupling the storage subsystems

Abstract Paragraph:
A system includes plural storage subsystems each having a controller and an expander to couple to storage devices. The controller accesses the storage devices through the expander, and the expander has interfaces for coupling to the storage devices. The system further includes an intercontroller link to connect expanders in two storage subsystems to enable the controller in one of the storage subsystems to communicate with the controller in another one of the storage subsystems through the expanders and the intercontroller link.

CLAIMS:

1. A system, comprising: plural storage subsystems, each storage subsystem having a controller, an expander, and zero or more storage devices coupled to the expander, the controller to access storage devices through the expander, and the expander having interfaces for coupling to storage devices; and an intercontroller link to connect expanders in different storage subsystems to enable the controller in one of the storage subsystems to communicate with the controller in another one of the storage subsystems through the expanders and the intercontroller link.

13. A method for use in a system having plural storage subsystems, each storage subsystem having a controller and an expander, the method comprising: accessing, by the controller in a first one of the storage subsystems, a storage device in the first storage subsystem through the expander in the first storage subsystem; and communicating over an intercontroller link that connects the expander in the first storage subsystem with an expander in a second one of the storage subsystems, wherein the controller in the first storage subsystem communicates with the controller in the second storage subsystem through the intercontroller link and the expanders in the first and second storage subsystems.

18. An expander in a first storage subsystem, comprising: a first interface to couple to a storage device; a second interface to couple to an intercontroller link to connect the expander in the first storage subsystem with an expander in a second storage subsystem; and a controller to communicate with another controller in the second storage subsystem through the second interface and the intercontroller link.

Previous Doc Next Doc Go to Doc#

First Hit · Previous Doc Next Doc Go to Doc#

Generate Collection Print

L4: Entry 1 of 4 File: PGPB Mar 24, 2005

PGPUB-DOCUMENT-NUMBER: 20050066100
PGPUB-FILING-TYPE: new
DOCUMENT-IDENTIFIER: US 20050066100 A1

TITLE: System having storage subsystems and a link coupling the storage subsystems

PUBLICATION-DATE: March 24, 2005

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
Elliott, Robert C.	Houston	TX	US	
Grieff, Thomas	Cypress	TX	US	
Foster, Joseph E.	Houston	TX	US	

APPL-NO: 10/ 669388 [PALM]
DATE FILED: September 24, 2003

INT-CL: [07] G06 F 13/00

US-CL-PUBLISHED: 710/300
US-CL-CURRENT: 710/300

REPRESENTATIVE-FIGURES: 1B

ABSTRACT:

A system includes plural storage subsystems each having a controller and an expander to couple to storage devices. The controller accesses the storage devices through the expander, and the expander has interfaces for coupling to the storage devices. The system further includes an intercontroller link to connect expanders in two storage subsystems to enable the controller in one of the storage subsystems to communicate with the controller in another one of the storage subsystems through the expanders and the intercontroller link.

Previous Doc Next Doc Go to Doc#

First Hit · Fwd Refs Previous Doc Next Doc Go to Doc#

End of Result Set

Generate Collection Print

L4: Entry 4 of 4 File: USPT Jan 11, 2005

DOCUMENT-IDENTIFIER: US 6842829 B1
TITLE: Method and apparatus to manage independent memory systems as a shared volume

Detailed Description Text (7):
In accordance with a preferred embodiment of the present invention, the switched architecture includes path 190 between switch 130 and switch 180. Path 190 is a switch-to-switch path that allows for inter-controller access to memory systems and input/output (I/O) interfaces in a redundant controller environment. For example, when a request is received on host CA 102, CPU 110 may access a device via drive CA 156 through path 190. As a further example, when a request is received on host CA 154, CPU 160 may access RAM 128 via RMC in 124 through path 190.

Previous Doc Next Doc Go to Doc#

First Hit · Fwd Refs Previous Doc Next Doc Go to Doc#
End of Result Set

Generate Collection Print

L4: Entry 4 of 4 File: USPT Jan 11, 2005

US-PAT-NO: 6842829
DOCUMENT-IDENTIFIER: US 6842829 B1

TITLE: Method and apparatus to manage independent memory systems as a shared volume

DATE-ISSUED: January 11, 2005

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Nichols; Charles F.	Wichita	KS		
Holt; Keith W.	Wichita	KS		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
LSI Logic Corporation	Milpitas	CA			02

APPL-NO: 10/ 006162 [PALM]
DATE FILED: December 6, 2001

INT-CL: [07] G06 F 12/00

US-CL-ISSUED: 711/147; 711/114, 714/6
US-CL-CURRENT: 711/147; 711/114, 714/6

FIELD-OF-SEARCH: 711/114, 711/147, 714/6

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	5606706	February 1997	Takamoto et al.	
☐	5640530	June 1997	Beardsley et al.	
☐	5742792	April 1998	Yanai et al.	711/162
☐	5895485	April 1999	Loechel et al.	
☐	6044444	March 2000	Ofek	
☐	6085333	July 2000	DeKoning et al.	714/7
☐	6101497	August 2000	Ofek	707/10
☐	6219751	April 2001	Hodges	711/114
☐	6321298	November 2001	Hubis	711/124

☐	6349357	February 2002	Chong, Jr.	711/111
☐	6442551	August 2002	Ofek	
☐	6557079	April 2003	Mason, Jr. et al.	

ART-UNIT: 2189

PRIMARY-EXAMINER: Sparks; Donald

ASSISTANT-EXAMINER: Chace; Christian P.

ATTY-AGENT-FIRM: Yee & Associates

ABSTRACT:

A switched architecture is provided to allow controllers to manage physically independent memory systems as a single, large memory system. The switched architecture includes a path between switches of controllers for inter-controller access to memory systems and input/output interfaces in a redundant controller environment. Controller memory systems are physically independent of each other; however, they are logically managed as a single, large memory pool. Cache coherency is concurrently maintained by both controllers through a shared locking mechanism. Volume Logical Block Address extents or individual cache blocks can be locked for either shared or exclusive access by either controller. There is no strict ownership model to determine data access. Access is managed by the controller in the pair that receives the access request. When a controller is removed or fails, a surviving controller may take appropriate action to invalidate all cache data that physically resides in the failed or missing controller's memory systems. Cached write data may be mirrored between redundant controllers to prevent a single point of failure with respect to unwritten cached write data.

9 Claims, 7 Drawing figures

Previous Doc Next Doc Go to Doc#

First Hit · Fwd Refs Previous Doc Next Doc Go to Doc#
End of Result Set

Generate Collection Print

L4: Entry 4 of 4 File: USPT Jan 11, 2005

US-PAT-NO: 6842829
DOCUMENT-IDENTIFIER: US 6842829 B1

TITLE: Method and apparatus to manage independent memory systems as a shared volume

DATE-ISSUED: January 11, 2005

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Nichols; Charles F.	Wichita	KS		
Holt; Keith W.	Wichita	KS		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
LSI Logic Corporation	Milpitas	CA			02

APPL-NO: 10/ 006162 [PALM]
DATE FILED: December 6, 2001

INT-CL: [07] G06 F 12/00

US-CL-ISSUED: 711/147; 711/114, 714/6
US-CL-CURRENT: 711/147; 711/114, 714/6

FIELD-OF-SEARCH: 711/114, 711/147, 714/6

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	5606706	February 1997	Takamoto et al.	
☐	5640530	June 1997	Beardsley et al.	
☐	5742792	April 1998	Yanai et al.	711/162
☐	5895485	April 1999	Loechel et al.	
☐	6044444	March 2000	Ofek	
☐	6085333	July 2000	DeKoning et al.	714/7
☐	6101497	August 2000	Ofek	707/10
☐	6219751	April 2001	Hodges	711/114
☐	6321298	November 2001	Hubis	711/124

☐	6349357	February 2002	Chong, Jr.	711/111
☐	6442551	August 2002	Ofek	
☐	6557079	April 2003	Mason, Jr. et al.	

ART-UNIT: 2189

PRIMARY-EXAMINER: Sparks; Donald

ASSISTANT-EXAMINER: Chace; Christian P.

ATTY-AGENT-FIRM: Yee & Associates

ABSTRACT:

A switched architecture is provided to allow controllers to manage physically independent memory systems as a single, large memory system. The switched architecture includes a path between switches of controllers for inter-controller access to memory systems and input/output interfaces in a redundant controller environment. Controller memory systems are physically independent of each other; however, they are logically managed as a single, large memory pool. Cache coherency is concurrently maintained by both controllers through a shared locking mechanism. Volume Logical Block Address extents or individual cache blocks can be locked for either shared or exclusive access by either controller. There is no strict ownership model to determine data access. Access is managed by the controller in the pair that receives the access request. When a controller is removed or fails, a surviving controller may take appropriate action to invalidate all cache data that physically resides in the failed or missing controller's memory systems. Cached write data may be mirrored between redundant controllers to prevent a single point of failure with respect to unwritten cached write data.

9 Claims, 7 Drawing figures

Previous Doc Next Doc Go to Doc#

L3: Entry 90 of 92 File: USPT Dec 13, 1988

DOCUMENT-IDENTIFIER: US 4791629 A
TITLE: Communications switching system

Brief Summary Text (14):
The system software consists of two levels. A lower level is a real time code which is interrupt driven and is imlemented in an assembly code appropriate to the processor of each controller. The real time code works synchronously to control eight digital or eight analog interfaces and intercontroller communications between controllers.

Detailed Description Text (7):
As noted previously, the present invention utilizes a unique two level software configuration in which a synchronous functional code is accessed as necessary without disturbing the operation of the synchronous real time code. The real time code and functional code communicate with each other via FIFOs and exchange information such as intercontroller packet information, station display information, key code information and trunk status information, etc.

Detailed Description Text (28):
1. The callable routine calls STNINIT or TMINIT depending on whether the current controller is connected to a digital interface or an analog interface. STNINIT initializes all data structures and variables in the processor 29 memory space pertaining to digital stations. TMINIT does the same for analog devices.

Detailed Description Text (40):
3. Outputting the intercontroller output packets and determining when to output the intercontroller packets so that all other controllers in the system have a chance to read them.

Detailed Description Text (41):
4. Inputting a new intercontroller input packet from a different controller.

Detailed Description Text (47):
Data values are stored into SE buffer memory which are read by the SEP 64 and stored into frame space. These values include the next intercontroller corresponding address, the link input and output channel pointer, and data values stored in the IEOFIFO. Values in the IEOFIFO can be any values that are in frame space including corresponding addresses, time slot registers, link control and codec enable etc.

Detailed Description Text (48):
3. INTERCONTROLLER OUTPUT PACKET PROCESS

Detailed Description Text (49):
This process determines whether the next output packet is a NIL packet, status packet, low speed packet, or high speed packet. If the next packet is a NIL packet, a NIL packet is stored into the IC packet output buffer in the SE memory. If the packet is a status packet, the STAT packet is moved to the SE. If a low speed packet is seen, the next packet in the PFIFO is moved into the SE. Otherwise, the next packet is a high speed packet and the next packet in the HS BJFF is stored. A high speed packet is a packet from the HS BJFF. A new packet is sent every superframe period. This kind of packet communication is used for downloading purposes. The current controller is locked onto the destination controller and the intercontroller corresponding address in frame space will not be changed to another controller port every superframe. A low speed packet is a packet POFIFO the same packet is sent every 32 milliseconds. Therefore, this packet is not removed from the top of the FIFO until 32 milliseconds are up. Therefore, it will be guaranteed that every controller in the system will

receive the packet.

Detailed Description Text (54):
6. INTERCONTROLLER PACKET CYCLE PROCESSOR

Detailed Description Text (55):
This process is invoked once every 32 milliseconds. The process removes the top packet on the PROFIFO and resets the HS BUFF output pointer to the beginning. It then determines whether it is time to transmit a status packet. A status packet is transmitted from the lOCAL FIFO every 256 milliseconds. Actual status information is looked up in the local directory. Next, the process determines whether there are any low speed packets into PROFIFO FIFO. Or any high speed packets in the HS BUFF. After the next packet is determined, the packet source flag (P source) is set to be used in the intercontroller output process (POPROC) to mark the FIFO with the next ready packet. The packet is then processed in the next interrupt cycle.

Detailed Description Text (100):
The prologue sets up correspondent address (port) for intercontroller reading. It also resets ESF (early superframe) addresses at frame 7 and late superframe (LSF) addresses at frame 1. Further, the prologue moves the IE register from superframe space to frame space.

Detailed Description Text (136):
The ten active time slots for each controller are arranged in the preferred embodiment as follows. Time slot 0 (of the 10 assigned time slots) carries an ICP (Inter Controller Packet) which is used to broadcast control messages to all controllers. Time slot contains RS-232 data and is used only by the administrative and maintenance controllers. No other controllers use this slot. Time slots 2-9 contain station voice data for digital controllers, for reading station voice data to another station. Analog controllers utilize slots 2-9 for reading trunk voice data to another trunk. Administrative and conferencer controllers utilize time slots 2-9 for conference information to input 2-4 way voice summers to allow conference calling.

Detailed Description Text (151):
The controller provides a means of controlling communication between the Xbus 23 and the digital stations or analog interfaces to which each controller is coupled. Each controller provides storage for incoming and outgoing data as well as switching services for the individual lines coupled to each controller. The controllers are shown in more detail in FIG. 2. Each controller includes a buffer and latch 41 coupled through line 43 to microprocessor 29. The processor 29 is coupled through an address decoder 69 to buffer 41 and to a memory consisting of dynamic RAM 40, static RAM 38, and EPROM 37. The processor 29 is also coupled through line 46 to a memory page control 39. The processor 29 is coupled through line 44 directly to the dynamic RAM 40, static RAM 38, and eprom 37. Coupled to the static RAM 38, is a standby power block 30 which provides battery backup during power failure. The clock receiver is coupled to the clock bus 22, switching element 32 and interface element 33. Power is supplied through a DC voltage regulator 34 and a power up/down circuit 31 provided on each controller sequences the activation and deactivation of the microprocessor 29, the SE 32, and the static RAM 38. The switching element 32 is coupled to th interface element 33. The interface element is coupled to the Xbus 23 through a transceiver 36. The interface element 33 is also coupled to digital station interfaces and to analog device interfaces through serial data bus 42, link/codec select line 70, and link control signal line 71.

Detailed Description Text (176):
The processor buffer consists of a double buffering scheme which is similar to the link/codec buffer and Xbus buffer. The buffer consists of two buffers 54A and 54B with each having an input portion and an output portion. These buffers are memory mapped into the address space of the processor 29 so that the real time code has access to one of the buffers at a particular time while the SEP 64 will process information in the other buffer. The processor 29 access and SEP 64 switch buffers every millisecond (superframe). This means that every 125 microseconds only one byte of link control information and one byte of intercontroller information is swapped. A total of one millisecond is needed to fill the link packet input buffer and intercontroller packet input buffer.

Detailed Description Text (178):
The 8 byte intercontroller packet consists of an 8 byte packet used for intercontroller communications. The 8 byte field names are as follows:

Detailed Description Text (187):
Intercommunication packets are passed between controllers as a way to pass data and status information to each other via the Xbus 23. The SEP 64 will move a byte of packet information from the intercontroller input port to the processor 29 intercontroller input buffer and from the intercontroller output buffer to the intercontroller output port every frame.

Detailed Description Text (191):
As noted previously, each controller has 16 ports. Port 0 is for intercontrolled communications and accepts intercontroller packets. All controllers utilize port 0. Port 1 is utilized only by digital controllers for voice override features.

Detailed Description Text (207):
SAM 68 is a special access memory which is used to configure the interface element. The SAM 68 is accessible to the switching element 32 via the SE interface 56 in the same way that the IE RAM 65 is accessible. The contents of the SAM 68 are available fulltime to the interface element module which uses the data. As can be seen in FIG. 7, the SAM address locations 75 are locations 224 through 231 of the IE memory 65. Address 224 is time slot base and address 225 is time slot count. The contents of these registers are used to determine the time slots which are assigned to an individual controller. The time slot counter is reset to 0 after each frame. If the time slot count equals the time slot base, that controller becomes active for the number of time slots indicated in the time slot count register. In the preferred embodiment, the number of assigned time slots is 10. Address 226 is the controller address which is 5 bits in length. Addresses 227 through 230 are mode data locations used by the link/codec interface module. These signals control the links to the digital stations and analog interfaces. Address 231 includes an uplink dataword which indicates uplink errors and must be cleared by the switching element 32 via the SE interface 56.

Detailed Description Text (224):
Controllers communicate with each other by intercontroller packets which are transmitted once each superframe. During each frame of a superframe, data is transmitted to the Xbus buffer 73 of the interface element of a controller. After each frame, the SEP (64) of the receiving controller accesses the Xbus buffer 73 and loads the data into the superframe buffer. In order to transfer all of the data in a single superframe, the superframe buffers 54A and 54B of the receiving controller are switched with an ESF signal after the initial frame of the superframe. After all eight frames of the data packet have been stored in th superframe buffer (54A for example) the superframe buffers are switched, making buffer 54A visible to the processor 29. The packet is an 8 byte packet laid out as follows. Two bytes are a destination field identifying the receiving controller. The next two bytes are a sender field identifying the transmitting controller. A one byte command field identifies the nature of the packet with the remaining 3 bytes a reference field containing supporting data.

Previous Doc Next Doc Go to Doc#

L3: Entry 90 of 92 File: USPT Dec 13, 1988

US-PAT-NO: 4791629
DOCUMENT-IDENTIFIER: US 4791629 A

TITLE: Communications switching system

DATE-ISSUED: December 13, 1988

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Burns; C. A.	Palo Alto	CA		
DeCoursey; Calvin H.	Reno	NV		
Junker; Hans H.	Mountain View	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
IBM Corporation	Santa Clara	CA			02

APPL-NO: 06/ 869580 [PALM]
DATE FILED: June 2, 1986

INT-CL: [04] H04J 3/16

US-CL-ISSUED: 370/85; 370/95
US-CL-CURRENT: 370/363; 370/458

FIELD-OF-SEARCH: 370/85, 370/95, 370/89, 370/86, 370/94, 370/110.1, 370/67

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4549291	October 1985	Re Noulin et al.	370/89
☐	4628504	December 1986	Brown	370/85
☐	4630263	December 1986	Townsend et al.	370/85
☐	4663758	May 1987	Lambarelli et al.	370/85

ART-UNIT: 263

PRIMARY-EXAMINER: Olms; Douglas W.

ATTY-AGENT-FIRM: Blakely, Sokoloff, Taylor & Zafman

ABSTRACT:

A communication switching system in which control functions are distributed homogeneously, as opposed to being distributed by individual function, throughout the system. A plurality of controllers are coupled to a time division multiplex (TDM) bidirectional bus. Each controller includes a microprocessor to perform control functions. An interface element on the controller accesses the TDM bus while a switching element on each controller controls access between the interface element and the microprocessor. The switching element and interface element include storage areas which are double buffered to allow for more efficient operation and use of the memory space. The TDM bus is divided into a plurality of time slots. Each controller is assigned certain of the time slots and utilizes those time slots to communicate with other controllers. The system includes provisions for dynamic allocation of time slots.

16 Claims, 16 Drawing figures

Previous Doc Next Doc Go to Doc#

L3: Entry 90 of 92 File: USPT Dec 13, 1988

US-PAT-NO: 4791629
DOCUMENT-IDENTIFIER: US 4791629 A

TITLE: Communications switching system

DATE-ISSUED: December 13, 1988

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Burns; C. A.	Palo Alto	CA		
DeCoursey; Calvin H.	Reno	NV		
Junker; Hans H.	Mountain View	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
IBM Corporation	Santa Clara	CA			02

APPL-NO: 06/ 869580 [PALM]
DATE FILED: June 2, 1986

INT-CL: [04] H04J 3/16

US-CL-ISSUED: 370/85; 370/95
US-CL-CURRENT: 370/363; 370/458

FIELD-OF-SEARCH: 370/85, 370/95, 370/89, 370/86, 370/94, 370/110.1, 370/67

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4549291	October 1985	Re Noulin et al.	370/89
☐	4628504	December 1986	Brown	370/85
☐	4630263	December 1986	Townsend et al.	370/85
☐	4663758	May 1987	Lambarelli et al.	370/85

ART-UNIT: 263

PRIMARY-EXAMINER: Olms; Douglas W.

ATTY-AGENT-FIRM: Blakely, Sokoloff, Taylor & Zafman

ABSTRACT:

A communication switching system in which control functions are distributed homogeneously, as opposed to being distributed by individual function, throughout the system. A plurality of controllers are coupled to a time division multiplex (TDM) bidirectional bus. Each controller includes a microprocessor to perform control functions. An interface element on the controller accesses the TDM bus while a switching element on each controller controls access between the interface element and the microprocessor. The switching element and interface element include storage areas which are double buffered to allow for more efficient operation and use of the memory space. The TDM bus is divided into a plurality of time slots. Each controller is assigned certain of the time slots and utilizes those time slots to communicate with other controllers. The system includes provisions for dynamic allocation of time slots.

16 Claims, 16 Drawing figures

Previous Doc Next Doc Go to Doc#

L3: Entry 87 of 92 File: USPT Apr 2, 1996

DOCUMENT-IDENTIFIER: US 5504926 A
TITLE: Method for a host central processor and its associated controller to capture the selected one of a number of memory units via path control commands

Brief Summary Text (4):
Workers are also familiar with the practice of associating one or several host central processors (CPUs) with an array of memory units via intermediate controllers (e.g. see U.S. Pat. Nos. 4,982,324, 4,413,317, 3,889,237, 4,183,084, 3,623,014). And, today, it is not uncommon to operatively associate an array of peripheral storage units with a host computer via an appropriate controller device. FIG. 1 schematically suggests this for a host H and a number of like memory units D, connectible to H via a controller unit C, including suitable interfaces 1F, 1F' (e.g. see ports 1-4, each connected to a respective, like-numbered drive-port).

Brief Summary Text (7):
Now, it may become desirable to mediate the competition between two host-computers to "capture" a selected memory unit (e.g. drive), using certain mediation means. Such a mediation means is the "Shared Memory Interface" (SMi) indicated in FIG. 3, shown in operative relation with eight dual-ported disk drives D0-D7 and a pair of host/controller arrays: host A16 with Controller A and host A17 with Controller B, plus suitable controller inter-faces, IPI-2, IPI-3. Workers should realize that mediation means SMi provides a communication link that allows one controller to request, or transmit, information to, or from, the other controller--here via what may be called "direct memory access" (DMA) something that requires special interfaces [e.g. additional special hardware, as "shared memory", SM, and associated interfaces I-A, I-B, (memory typically a dedicated random access memory RAM) and special software]. Shared memory, SM, also functions as a common storage unit that can be used to maintain "data path control": i.e. allow/disallow a host to write/read information to/from a selected disk drive.

Brief Summary Text (14):
A salient feature of this "Natural Linking" is that it operates "Controller-to-Controller", using only a disk drive as its link, i.e. linking via means which are already part of the system (e.g. here only via a disk drive normally coupled to a M-9730 Controller, plus the "standard" IPI-2 interface normally used by this Controller). That is, here, the inter-controller linkage is "natural", using an already-present disk drive unit via an IPI-2 interface which is normally used for disk control or read/write commands. This way, one controller can now, use its normal (natural) interface to communicate with the other controller,--both controllers being coupled to the two ports of each disk drive. This novel natural linkage allows a disk control unit (controller) to request or transmit information to, or from, its companion controller via a selected disk drive in the already-present array of drives. Several useful, surprising advantages will be evident to workers (e.g. no special interfaces needed, only normal "IPI-2" interface as in FIG. 4).

Brief Summary Text (15):
And, unlike the aforedescribed "Shared Memory Linkage", so using a disk drive to transmit path control information, and to interrupt the other controller for updates, eliminates the need for dedicated RAM or other shared memory to maintain "data path control" and for other associated hardware and software. Workers will recognize that this direct, natural intercontroller linkage gives each controller a means of maintaining a resident copy of path-control information, while still maintaining a modicum of data security (at the subsystem). Further, only a minor part of disk-storage space is needed (e.g. I have found that a mere three sectors of track space on the maintenance cylinder can store all requisite data-path control information).

First Hit · Fwd Refs Previous Doc Next Doc Go to Doc#

Generate Collection Print

L3: Entry 87 of 92 File: USPT Apr 2, 1996

US-PAT-NO: 5504926
DOCUMENT-IDENTIFIER: US 5504926 A

TITLE: Method for a host central processor and its associated controller to capture the selected one of a number of memory units via path control commands

DATE-ISSUED: April 2, 1996

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Jackson; Gary E.	Lake Forest	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Unisys Corporation	Blue Bell	PA			02

APPL-NO: 07/ 949967 [PALM]
DATE FILED: September 24, 1992

INT-CL: [06] G06 F 12/00, G06 F 12/02, G06 F 13/00, G06 F 13/14

US-CL-ISSUED: 395/825; 395/856, 395/475, 395/439, 364/131, 364/228.3, 364/245.5, 364/245.7
US-CL-CURRENT: 710/5; 700/2, 710/36, 711/112, 711/148

FIELD-OF-SEARCH: 364/200, 364/131, 395/275, 395/325, 395/200, 395/200.01, 395/425, 395/725, 395/825, 395/856, 395/475, 395/439

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4455621	June 1984	Pelley et al.	364/900
☐	4747043	May 1988	Rodman	364/200
☐	4759010	July 1988	Murata et al.	370/58
☐	5060140	October 1991	Brown et al.	364/200
☐	5261072	November 1993	Siegel	395/425
☐	5295247	March 1994	Chang et al.	395/325
☐	5325488	June 1994	Carteau et al.	395/275

ART-UNIT: 237

PRIMARY-EXAMINER: Lee; Thomas C.

ASSISTANT-EXAMINER: Krick; Rehana

ATTY-AGENT-FIRM: McCormack; John J. Starr; Mark T.

ABSTRACT:

A pair of data processing systems, each of the data processing system having a host central processor and an associated controller including memory, both of the data processing systems to be cooperatively associated with a number of disk drive memory units, each of the disk drive memory unit coupled to both said controllers. Either one of the host central processors can appropriate any one of the disk drive memory units as a selected disk drive memory unit by propagating path-control-data to the memory in both of the controllers arid in the selected disk drive memory unit.

3 Claims, 11 Drawing figures

Previous Doc Next Doc Go to Doc#

L3: Entry 86 of 92 File: USPT Jun 25, 1996

DOCUMENT-IDENTIFIER: US 5530845 A

**** See image for Certificate of Correction ****

TITLE: Storage control subsystem implemented with an application program on a computer

Abstract Text (1):

A storage controller is disclosed which may emulate several types of specialized host specific and/or storage device specific storage controllers. The storage controlling system can transfer information between one or more different types of target units and one or more channels of at least one host. The system is provided with a computer, which includes a first interface, a second interface, and a programmable storage controller. The first interface is configured to receive one or more channel adapters which carry one or more channel programs transmitted from the channels of the host. The channel programs may carry data, status information, and commands. The second interface allows input and output to storage facilities which comprise one or more target units. The programmable storage controller may be provided with a device coupled to the channel adapters for translating channel program commands, and determining, from the channel program, a target unit for which at least one channel program is transmitted. A set of equipment controllers is provided which interpret channel program commands and status information, and which further control data transfers to and from the storage facilities in accordance with the channel program command. A device is also provided for establishing a unit thread by choosing an equipment controller from the set of equipment controllers as a function of the type of equipment that the channel requests as a target.

Brief Summary Text (6):

At the present time, due to rapid advances in peripheral technologies, newly developed storage facilities are available which have increased capabilities in areas such as efficiency and size. For example, rewritable optical disks, optical tapes, and 4 mm digital tapes (DAT) are each known for their large storage capacities, reasonably quick access time, and low floor space and power requirements. However, in order for current mainframe and mini computers to access these improved storage facilities, specialized storage controllers must be developed (or purchased) which may handle such storage facilities. For example, in order to facilitate connection of mainframe FIPS 60 channels to a SCSI storage device interface, several specialty manufacturers provide plug-in boards which allow VME base computers, such as the SUN, to intercept FIPS 60 channel inputs. However, these products have been provided by small organizations for limited specialty applications; they are not generic storage controllers which will support all or a significant portion of newly developed storage facilities.

Brief Summary Text (12):

It is yet a further object of the present invention is to provide a storage controlling system which has an intercontroller communication bridge, which allows controller facilities to be shared.

Brief Summary Text (20):

The present invention, therefore, is directed to a storage controlling system for transferring information between one or more target units and one or more channels of at least one host. The host is configured for one or more types of equipment corresponding to the one or more target units. The system is provided with a computer, which includes a first interface, a second interface, and a programmable storage controller. The first interface is configured to receive one or more channel adapters which carry one or more channel programs transmitted from the one or more channels of the host. Each channel program typically carries data, status information and commands. The second interface interfaces, and thus allows input and output, to storage facilities which comprise the one or more target units. In a particular aspect, the type of equipment for which the host is configured is different than the one or more target units.

Brief Summary Text (22):
In accordance with another aspect of the present invention, the equipment controller is further provided with a mechanism for calling a cache manager. In addition, the equipment controller may be provided with a device for prioritizing input and output to and from the storage facilities. The one or more channel adapters may comprise one or more channel interface circuits, and the first interface may be provided with a channel interface controller. In this regard, the second interface may be provided with at least one of random access and sequential storage device channel adapters.

Brief Summary Text (24):
In yet another aspect of the present invention, the programmable storage controller is further provided with a dispatcher for controlling the operation of the programmable storage controller. The channel interface controller is provided with an interrupt processor, responsive to a command received by the one or more channel programs, for interrupting the dispatcher, thus causing establishment of another unit thread, and a device for calling the execution of the another unit thread. The channel interface controller may also be provided with a device for controlling the one or more channel interface circuits and further a device for retrieving at least one of the channel programs.

Brief Summary Text (30):
In yet a further aspect of the present invention, the second interface is further provided with an interface to another storage controller, wherein the programmable storage controller comprises a communication bridge for communicating with the another storage controller. The programmable storage controller may be provided with a plurality of customized controller services, such as a caching algorithm, and a device for performing data compression and decompression.

Detailed Description Text (9):
Referring to FIG. 2, programmable storage controller 14 is provided with a channel adapter interface (CI) 210, an application interface (or dispatcher) 212, one or more controller emulators (equipment controllers) 214, a storage control manager 216, a cache manager 218, an I/O facilities handler (IOF) 220, a set of compression/decompression routines 222, and a device driver 224. Channel adapter interface 210, which is preferably implemented with software, provides a mechanism which allows application interface 212 to control other portions of the programmable storage controller, and also control channel adapter boards 26 (and thus the storage controller channels 13 of host system 12). CI 210 provides data, status, and commands from the channel, in the form of a channel program, to one or more controller emulators 214 which control data and status command exchanges with one or more target units. Application interface 212 handles command requests which are detected by CI 210 as incoming data streams. CI 210 is provided with an interrupt processor which causes execution thereof, and thus sets application interface 212 into action.

Detailed Description Text (10):
Application interface 212 routes the incoming commands (in the form of channel programs) to one or more controller emulators 214, by establishing and scheduling one or more unit threads based on the particular device (target unit) requested by the host system through the channel program. Application interface 212 uses a configuration table to determine which controller emulator should receive a request based upon the type of equipment the channel expects as a target. If the controller emulator which is needed to establish the unit thread is not currently loaded in memory of computer 15, application interface 212 will request the program to be loaded.

Detailed Description Text (17):
Host system 12 physically attaches to a storage controller through channel 13 which uses a unique addressing scheme. The programmable storage controller 14 may communicate with multiple channel interface cards of the channels, thus allowing multiple host channels to communicate with a single storage subsystem (i.e., the combined system of programmable storage controller 14 and storage facility 16). This allows multiple assembly host architectures (e.g., the IBM 370 Class, and Unisys 2200 Class) to communicate with a single parameter storage controller. The adapter card which is in the channel of the host system receives and passes on channel program commands and requests which are configured in a channel specific format. Meanwhile, application interface 212, which knows what type of host is connected to the channel adapter interface of the host, chooses an appropriate controller emulator which corresponds to the host, and assigns a unit thread representing that controller emulator.

Detailed Description Text (19):
A particular embodiment of programmable storage controller 14 will now described with reference to FIGS. 3a-3d. The application interface (or dispatcher) flow of operation is described in FIG. 3b. During execution of the application interface, in step S21, a packet is built into interface 18 for the CI which contains subchannel and command table information. Thereafter, in step S22, execution of the CI initialization (FIG. 3a, steps S1-S5) is called. Upon completion of CI initialization, which results in the execution of steps S1-S5 (FIG. 3a), the application interface then builds one or more unit thread semaphores; one unit thread semaphore is built for each device configured in the host system. Thereafter, in step S24, the one or more unit threads are scheduled, and thus initialized. Upon completion of scheduling/initialization of the unit threads, in step S25, a "get command" semaphore, is created. Then, steps S26-S28 are executed in response to clearance of the "get command" semaphore as indicated at step S26, which queues on the "get command" semaphore. Upon building an initial selection request packet for the CI in step S27, the application interface calls the CI for execution of the read channel processing (steps S6-S15, FIG. 3a).

Detailed Description Text (21):
The CI (i.e., the channel adapter interface) functions as a device driver for carrying out I/O to and from the one or more storage controller channels 13 of host system 12. Thus, whenever data is either written to or read from the one or more storage controller channels 13, the CI must be executed in some fashion. In order to execute a "read channel" function, the read channel portion of the CI is called (see steps S6-S15). In order to execute a "write channel" function, the "write to" channel portion of the CI is called for execution (steps S16-S20).

Detailed Description Text (24):
In step S35, the stored data which must be referenced is so referenced. That is, the storage/cache manager finds the data and returns an index to the controller emulator indicating where the controller emulator may find the data. Once this data is indexed in step S35, the controller emulator then makes another determination at step S36 as to what type of command is to be executed as instructed by the storage controller channel of the host system. If it is a read command, the controller emulator sends the data and status to the host by calling the CI (write to channel) (steps S16-S20), and returns to step S30 where it sleeps until the unit thread semaphore is again cleared. If the command by the host is a write command, the controller puts the data to be written in the cache buffer and sends a status back to the host by the use of the CI (write to channel) program. Upon sending this status, the controller emulator returns to step S30. After execution of each of steps S37, S38, S39, and/or S40, prior to returning to step S30, in step 40.1, the controller emulator clears the get command semaphore, thus allowing the application interface to execute steps S26-S38 and retrieve the appropriate command for use by the controller emulator.

Detailed Description Text (26):
The read channel processing of the CI comprises steps S6-S15. In step S6, the interrupts of the storage controller channels are disabled. Upon disabling of the interrupts of the storage controller channel interface card, in step S7, the CI sets the initial selection function, which in turn allows the host to perform initial device selection. Once the initial selection function is set, the CI waits for a device (target unit) selection by the storage controller channel in step S8. Then, in step S9, the CI puts the address and command of the device (target unit) selection in the initial selection request packet, which is built into step S27 of the application interface (FIG. 3b). Thereafter, in step S10, the unit thread semaphore corresponding to a desired controller emulator is cleared, thus causing the appropriate controller emulator to be executed. In step S11, the CI waits for the read data semaphore to be cleared. This occurs in step S31 of the controller emulator, where the controller emulator gets the command and parameters from the CI. If an I/O is necessary, in step S12, the CI performs such channel I/O. Thereafter, in step S13, the data requested by the host channel (subsystem data) is sent to the host channel. The read complete semaphore is then cleared in step S14, and the processing is returned to where it was called, which is, in this case, step S28 of the application interface.

Detailed Description Text (31):
In calling of either of the channel adapter interface 210, and the device driver 224 an IOTCL interface is used to avoid the overhead of the host controller operating system file structure system checks. This allows more efficient access directly to the media of the storage facilities.

Detailed Description Text (37):
FIG. 4 illustrates a block diagram of a further embodiment of the present invention. In this embodiment, host system 12 is provided with first and second host computers 28, 30, and an additional "other" storage controller 32 is connected to interface 20 of computer 15. In order to effect communication between programmable storage controller 14 and other storage controller 32, a mechanism must be provided for interfacing other storage controller 32 with appropriate portions of programmable storage controller 14, so as to allow communication of various commands, status and data therebetween. This mechanism may include an inter-controller communication bridge. The inter-controller communication bridge may comprise of a separate controller emulator which provides translation functions allowing the host computer to communicate with storage controller 32. The controller emulator makes the appropriate translation of host requests to the generic format, which is then interpreted by the storage control manager. The storage control manager then further translates the instructions and commands into an appropriate protocol which may be interpreted by the other storage controller 32 as actions or commands to respond to.

Detailed Description Text (38):
This additional inter-controller communication bridge may be provided in order to accommodate functions such as backing up one storage channel with another, or to provide alternate routes of communication in case a channel interface or storage controller has been disabled because of hardware or software failure. It is noted that although FIG. 4 indicates a particular configuration of the inter-controller connection, other storage controller 32 may be connected to first interface 18 rather than second interface 20.

Detailed Description Text (39):
In order for the programmable storage controller 14 to accommodate a number of different host computers, such as first host computer 28 and second host computer 30, each unit thread which is established and scheduled by application interface 212 (see FIG. 2) should further be provided with an interpretor which may interpret the different channel languages and/or channel dialects of the communication interface used by the one or more host computers.

Detailed Description Text (42):
Another embodiment of the storage controlling system of the present invention is depicted in FIG. 5. As shown in this figure, one or more mainframe computers, which may be, for example, IBM mainframes, are connected to mainframe interface hardware 52 of a microcomputer 54. Microcomputer 54 includes an IBM compatible personal computer with, for example, either an 80486 or 80586 microprocessor. Microcomputer 54 is configured with a programmable storage controller 14 having a SCSI device driver 56, which interfaces with a plurality of SCSI interfaces connected to optical disks 60, which comprise the storage facilities in the present embodiment.

CLAIMS:

1. A storage control subsystem connected between one or more storage controller channels of at least one host system and data storage facilities comprising a plurality of target units, said storage control subsystem comprising:

a programmable storage controller that emulates a plurality of types of target unit specific storage controllers, said programmable storage controller being implemented with an application program and a computer, said computer being configured by said application program;

a first interface for interfacing a plurality of channel adapters which carry a plurality of channel programs transmitted from the channels of the host system to said programmable storage controller, each channel program having means for carrying data, status information and commands; and

a second interface for interfacing said programmable storage controller to said target units;

said programmable storage controller comprising a plurality of controller emulators, said controller emulators comprising means for translating said channel programs and commands from a channel specific format to a generic format of said programmable storage controller that includes generic address and request information, to thereby facilitate data and status command exchanges with said plurality of target units.

11. The storage control subsystem according to claim 5, wherein said programmable storage controller further comprises a dispatcher for controlling the operation of said programmable storage controller, and further wherein said channel interface controller comprises an interrupt processor, responsive to a command received by each of said channel programs, for interrupting said dispatcher, thus causing said means for executing to execute another unit thread.

45. A storage control subsystem connected between storage controller channels of at least one host system and data storage facilities comprising a plurality of target units, said storage control subsystem comprising:

a programmable storage controller that emulates a plurality of types of target unit specific storage controllers, said programmable storage controller being implemented with a general purpose computer with a general purpose operating system supporting an application program, said general purpose computer being configured by said application program;

a first interface for interfacing a plurality of channel adapters which carry a plurality of channel programs transmitted from the channels of the host system to said programmable storage controller, each channel program having means for carrying data, status information and commands; and

a second interface for interfacing said programmable storage controller to said target units;

said programmable storage controller comprising a plurality of controller emulators, said controller emulators comprising means for translating said channel programs and commands from a channel specific format to a generic format of said programmable storage controller that includes generic address and request information, to thereby facilitate data and status command exchanges with said plurality of target units.

Previous Doc Next Doc Go to Doc#

L3: Entry 86 of 92 File: USPT Jun 25, 1996

US-PAT-NO: 5530845
DOCUMENT-IDENTIFIER: US 5530845 A
** See image for Certificate of Correction **

TITLE: Storage control subsystem implemented with an application program on a computer

DATE-ISSUED: June 25, 1996

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Hiatt; David M.	St. Louis	MO		
Klos; Timothy R.	St. Louis	MO		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Southwestern Bell Technology Resources, Inc.	St. Louis	MO			02

APPL-NO: 08/ 373896 [PALM]
DATE FILED: January 17, 1995

PARENT-CASE:
This application is a continuation of application Ser. No. 07/882,010, filed May 13, 1992, now abandoned.

INT-CL: [06] G06 F 13/10

US-CL-ISSUED: 395/500; 395/882, 395/883
US-CL-CURRENT: 703/27; 710/62, 710/63

FIELD-OF-SEARCH: 395/500, 395/882, 395/883

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	3544969	December 1970	Rakoczi et al.	
☐	4084235	April 1978	Hirtle et al.	
☐	4277827	July 1981	Carlson et al.	
☐	4394732	July 1983	Swenson	
☐	4394733	July 1983	Swenson	
☐	4399503	August 1983	Hawley	
☐	4415969	November 1983	Bayliss et al.	

☐	4425615	January 1984	Swenson et al.	
☐	4454595	June 1984	Cage	
☐	4467421	August 1984	White	
☐	4476526	October 1984	Dodd	
☐	4500933	February 1985	Chan	
☐	4511963	April 1985	Kantner	
☐	4638423	January 1987	Ballard	
☐	4775969	October 1988	Osterlund	
☐	4792896	December 1988	MacLean et al.	
☐	4803623	February 1989	Klashka et al.	395/275
☐	4812975	March 1989	Adachi et al.	
☐	4855905	August 1989	Estrada et al.	
☐	4864291	September 1989	Korpi	
☐	4864532	September 1989	Reeve et al.	
☐	4868734	September 1989	Idleman et al.	
☐	4888691	December 1989	George et al.	
☐	4888727	December 1989	Getson, Jr. et al.	
☐	4965801	October 1990	Dulac	
☐	5088033	February 1992	Binkley et al.	
☐	5131082	July 1992	Bonevento et al.	
☐	5151985	September 1992	Sander et al.	
☐	5247638	September 1993	O'Brien et al.	395/425

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO	PUBN-DATE	COUNTRY	US-CL
246125	March 1994	AR	

OTHER PUBLICATIONS

Copy of a diagram, for an adaptor to link an ECKD/byte multiplexer channel to Micro Channel.

ART-UNIT: 235

PRIMARY-EXAMINER: Lim; Krisna

ATTY-AGENT-FIRM: Greenblum & Bernstein

ABSTRACT:

A storage controller is disclosed which may emulate several types of specialized host specific and/or storage device specific storage controllers. The storage controlling system can transfer information between one or more different types of target units and one or more channels of at least one host. The system is provided with a computer, which includes a first interface, a second interface, and a programmable storage controller. The first interface is configured to receive one or more channel adapters which carry one or more channel programs transmitted from the channels of the host. The channel programs may carry data, status information, and commands. The second interface allows input and output to storage facilities which comprise one

or more target units. The programmable storage controller may be provided with a device coupled to the channel adapters for translating channel program commands, and determining, from the channel program, a target unit for which at least one channel program is transmitted. A set of equipment controllers is provided which interpret channel program commands and status information, and which further control data transfers to and from the storage facilities in accordance with the channel program command. A device is also provided for establishing a unit thread by choosing an equipment controller from the set of equipment controllers as a function of the type of equipment that the channel requests as a target.

48 Claims, 8 Drawing figures

Previous Doc Next Doc Go to Doc#

L3: Entry 83 of 92 File: USPT May 13, 1997

DOCUMENT-IDENTIFIER: US 5630169 A
TITLE: Apparatus for a host central processor with associated controller to capture a selected one of a number of memory units via path control commands

Brief Summary Text (4):
Workers are also familiar with the practice of associating one or several host central processors (CPUs) with an array of memory units via intermediate controllers (e.g. see U.S. Pat. Nos. 4,982,324, 4,413,317, 3,889,237, 4,183,084, 3,623,014). And, today, it is not uncommon to operatively associate an array of peripheral storage units with a host computer via an appropriate controller device. FIG. 1 schematically suggests this for a host H and a number of like memory units D, connectible to H via a controller unit C, including suitable interfaces 1F, 1F' (e.g. see ports 1-4, each connected to a respective, like-numbered drive-port).

Brief Summary Text (7):
Now, it may become desirable to mediate the competition between two host-computers to "capture" a selected memory unit (e.g. drive), using certain mediation means. Such a mediation means is the "Shared Memory Interface" (SMi) indicated in FIG. 3, shown in operative relation with eight dual-ported disk drives D0-D7 and a pair of host/controller arrays: host A16 with Controller A and host A17 with Controller B, plus suitable controller inter-faces, IPI-2, IPI-3. Workers should realize that mediation means SMi provides a communication link that allows one controller to request, or transmit, information to, or from, the other controller--here via what may be called "direct memory access" (DMA) something that requires special interfaces [e.g. additional special hardware, as "shared memory", SM, and associated interfaces I-A, I-B, (memory typically a dedicated random access memory RAM) and special software]. Shared memory, SM, also functions as a common storage unit that can be used to maintain "data path control": i.e. allow/disallow a host to write/read information to/from a selected disk drive.

Brief Summary Text (14):
A salient feature of this "Natural Linking" is that it operates "Controller-to-Controller", using only a disk drive as its link, i.e. linking via means which are already part of the system (e.g. here only via a disk drive normally coupled to a M-9730 Controller, plus the "standard" IPI-2 interface normally used by this Controller). That is, here, the inter-controller linkage is "natural", using an already-present disk drive unit via an IPI-2 interface which is normally used for disk control or read/write commands. This way, one controller can now, use its normal (natural) interface to communicate with the other controller,--both controllers being coupled to the two ports of each disk drive. This novel natural linkage allows a disk control unit (controller) to request or transmit information to, or from, its companion controller via a selected disk drive in the already-present array of drives. Several useful, surprising advantages will be evident to workers (e.g. no special interfaces needed, only normal "IPI-2" interface as in FIG. 4).

Brief Summary Text (15):
And, unlike the aforedescribed "Shared Memory Linkage", so using a disk drive to transmit path control information, and to interrupt the other controller for updates, eliminates the need for dedicated RAM or other shared memory to maintain "data path control" and for other associated hardware and software. Workers will recognize that this direct, natural intercontroller linkage gives each controller a means of maintaining a resident copy of path-control information, while still maintaining a modicum of data security (at the subsystem). Further, only a minor part of disk-storage space is needed (e.g. I have found that a mere three sectors of track space on the maintenance cylinder can store all requisite data-path control information).

First Hit · Fwd Refs Previous Doc Next Doc Go to Doc#

Generate Collection Print

L3: Entry 83 of 92 File: USPT May 13, 1997

US-PAT-NO: 5630169
DOCUMENT-IDENTIFIER: US 5630169 A

TITLE: Apparatus for a host central processor with associated controller to capture a selected one of a number of memory units via path control commands

DATE-ISSUED: May 13, 1997

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Jackson; Gary E.	Lake Forest	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Unisys Corporation	Blue Bell	PA			02

APPL-NO: 08/ 625662 [PALM]
DATE FILED: March 29, 1996

PARENT-CASE:
This is a division of U.S. Ser. No. 07/949,967, filed Sep. 24, 1992, now U.S. Pat. No. 5,504,926.

INT-CL: [06] G06 F 12/02, G06 F 13/14

US-CL-ISSUED: 395/825; 395/856, 395/475, 395/439, 395/200.01
US-CL-CURRENT: 710/5; 710/36, 711/112

FIELD-OF-SEARCH: 395/825, 395/856, 395/475, 395/200.01, 395/725, 395/439

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4455621	June 1984	Pelley et al.	364/900
☐	4747043	May 1988	Rodman	364/200
☐	4759010	July 1988	Murata et al.	370/58
☐	5060140	October 1991	Brown et al.	364/200
☐	5261072	November 1993	Siegel	395/425
☐	5295247	March 1994	Chang et al.	395/325
☐	5325488	June 1994	Carteau et al.	395/275

ART-UNIT: 237

PRIMARY-EXAMINER: Lee; Thomas C.

ASSISTANT-EXAMINER: Krick; Rehana Perveen

ATTY-AGENT-FIRM: McCormack; John J. Starr; Mark T.

ABSTRACT:

A pair of data processing systems, each of the data processing system having a host central processor and an associated controller including memory, both of the data processing systems to be cooperatively associated with a number of disk drive memory units, each of the disk drive memory unit coupled to both said controllers. Either one of the host central processors can appropriate any one of the disk drive memory units as a selected disk drive memory unit by propagating path-control-data to the memory in both of the controllers and in the selected disk drive memory unit.

3 Claims, 11 Drawing figures

Previous Doc Next Doc Go to Doc#

L3: Entry 80 of 92 File: USPT Aug 4, 1998

DOCUMENT-IDENTIFIER: US 5790775 A
TITLE: Host transparent storage controller failover/failback of SCSI targets and associated units

Brief Summary Text (11):
However, the direction taken by Idleman requires a multi-level storage controller implementation. Each controller in the dual-redundant pair includes a two-level hierarchy of controllers. When the first level or host-interface controller of the first controller detects the failure of the second level or device interface controller of the second controller, it re-configures the data path such that the data is directed to the functioning second level controller of the second controller. In conjunction, a switching circuit re-configures the controller-device interconnections, thereby permitting the host to access the storage devices originally connected to the failed second level controller through the operating second level controller of the second controller. Thus, the presence of the first level controllers serves to isolate the host computer from the failover operation, but this isolation is obtained at added controller cost and complexity.

Brief Summary Text (13):
Related copending U.S. application Ser. No. 08/071,710, in the name of Sicola et al., in contrast to Idleman, describes a dual-active, redundant storage controller configuration in which each storage controller communicates directly with the host and its own attached devices, the access of which is shared with the other controller. Thus, a failover operation may be executed by one of the storage controller without the assistance of an intermediary controller and without the physical reconfiguration of the data path at the device interface. Not specifically addressed is the nature of the host's involvement in such an operation.

Drawing Description Text (6):
FIG. 4 is a simplified block diagram of an individual one of the storage controllers shown in the dual-active, redundant pair of FIG. 1. FIG. 5 is a block diagram of the dual-active, redundant pair of storage controllers, each as individually depicted in FIG. 4, further illustrating the intercontroller communication therebetween.

Detailed Description Text (2):
With reference to FIG. 1, there is illustrated a computer system 10 having a host CPU 12 coupled to storage controllers 14, 14' (referred to singularly as simply "14") via a host interconnect bus 16. In the preferred embodiment, the host interconnect bus 16 is a "host side" SCSI bus and the storage controllers 14, 14', each having a host port or host interface 18 (shown as a SCSI host port) connected to the same SCSI bus (i.e., the host side SCSI bus 16), operate as a dual-active, redundant pair in accordance with the present invention. Because the SCSI bus is an industry standard bus, the host CPU can be any SCSI-compliant host computer. Connected to each controller is a maintenance terminal 20, a user interface which allows for user directed controller management and status reporting. Further included in the computer system 10 is a communication link 22, which provides a communications path between the controllers 14, 14'. Thus, the controllers 14, 14' use the communication link 22 to share configuration and status information. Caches 24, 24' are attached to the controllers 14, 14' for performance purposes, cache 24 being used by controller 14 and cache 24' being used by controller 14', respectively. Although each cache is primarily associated with only one of the controllers as thus described, it may be accessed directly by the other controller under certain circumstances. The host CPU 12 communicates through each storage controller 14 to one or more "device side" SCSI buses 26. Each device side SCSI bus 26 is connected to a different one of device ports 28 in a device interface 30 on each of the controllers.

Detailed Description Text (5):

Controller storage addressing will now be discussed with reference to FIG. 2. Controller Port Target LUN (PTL) addressing is the process by which the controller selects storage space within a specific physical storage device. This type of addressing can be illustrated by the structure shown in FIG. 2. The SCSI bus 26 attached to the device port 28 connects the controller 14 to the SCSI-2 devices 34 of the physical storage media 32. This SCSI bus may be referred to as the "device side" SCSI bus, as already indicated, or the "back end" SCSI bus or device port. In the preferred embodiment, the maximum number of the device ports 28 supported in the device interface 30 is 6. Thus, the device ports 28 are numbered 1 through 6 as shown. The I/O devices 34, numbered 0 through 6, correspond to device side SCSI IDs 0 through 6, respectively. In the preferred embodiment, the controller uses device side SCSI IDs 6 or 7. This assignment is dependent upon the physical slot location of the housing in which the controller is installed. In a dual-redundant controller configuration, one controller has a device side ID of 6 and the other a device side ID of 7. These device side IDs are used to communicate between the controller and the devices on the device side SCSI bus. Each ID identifies a unique device on one of the device side buses 26.

Detailed Description Text (7):
Host storage addressing for computer system 10 is illustrated in FIG. 3. A typical host device interface 36 in host CPU 12 includes host ports 38, numbered 1 through N. Each host port 38 in the host device interface 36 is connected to one of host side SCSI buses 16--also known as "front end" buses--containing devices. From the perspective of the host CPU 12, the controller 14 is one of these devices. The controller is connected to the host side SCSI bus 16 at its host port or host interface 18 and has a corresponding host side SCSI ID (also known as a SCSI "target" ID), and contains up to eight devices. Thus, to support certain high-level storage subsystem functions, such as RAID, the controller 14 presents the entire physical device configuration (from the controller device interface 30 down to the physical device level) to the host CPU 12 as a group of host logical units (LUNs) 40. The host logical unit 40 usually consists of storage data being distributed throughout more than one physical device (as with a RAIDset). The controller presents these logical units to the host as individually addressable, virtual devices. The user configures the host logical units. Controller LUNs and host LUNs may represent the same structure, but this is often not the case, since RAIDsets (as well as stripesets and mirrorsets) cannot be configured in such a one-to-one relationship.

Detailed Description Text (9):
The host CPU 12 thus sees and addresses the physical storage media through the host LUNs 40. Although they share the same name, controller LUNs and host LUNs are logical addresses for two different storage structures. Controller LUNs exist only on the controller's device interface 30, and SCSI host LUNs exist only on a SCSI host's device interface 36. They may represent the same structure, but only if the user configures up to 8 controller devices in a one-to-one relationship with the host. This situation rarely occurs under normal operations.

Detailed Description Text (10):
Therefore, the host storage addressing is the process by which a SCSI host CPU selects a host logical unit made up of physical devices connected to a SCSI controller. First, if the host is connected to more than one adapter (not shown as a separate functional component in FIG. 3), the host selects an adapter with one or more SCSI-2 buses. Second, the host CPU selects the host side SCSI-2 bus that has the controller connected to it, if there are two or more SCSI-2 host interfaces. Third, the host CPU selects one of the controller's host side SCSI target IDs on that port (bus). Lastly, the host CPU addresses the controller with the LUN of the desired host logical unit. Thus, the target ID and host LUN are specified by the host CPU in accessing a "unit". The controller translates the host LUN into the physical addresses required to allow the host access to the virtual device represented by the host LUN.

Detailed Description Text (13):
Further, the user interface/configuration manager 60 enables the user to set and review controller configuration information, such as controller ID. To place the controller in a dual-active, redundant controller configuration, the configuration parameters must be set up at initialization time. Configuration information is entered by the user at the maintenance terminal 20 (from FIG. 1) connected to the maintenance terminal port 50 and stored in the nonvolatile memory 46. It includes the total number of SCSI target IDs. Up to four unique SCSI target IDs may be specified in any combination between the controllers for a dual-active, redundant controller configuration. This allows support for up to 32 logical units (LUNs), 8 host LUNs per target. To set the preferred paths to balance the load and improve performance, preferred path qualifiers to define which target IDs are assigned to each controller. For equal

sharing of resources, each controller is assigned two unique and "preferred" IDs; however, the invention is not so limited. If preferred IDs are not specified for one of the controllers, that controller will not respond to any target ID on the host's SCSI bus. However, in a dual-redundant configuration, if the controller with set preferred IDs fails, the controller with no preferred IDs will still take over the targets of the failed controller.

Detailed Description Text (15):
Resuming the storage controller description and still referring to FIG. 4, the cache 24 (from FIG. 1) is shown in dashed lines, as it is implemented as a separate module in the preferred embodiment. Alternately, it could be located on the same module as the controller. Thus, it can be considered part of the controller function, although it is located on a separate module. Moreover, the controller need not include a cache in order to take advantage of the present invention. Most modern storage controllers do incorporate a cache to enhance memory performance. If cache is employed, its design may utilize either a write-through or write-back caching policy. The separate buses containing the shared memory 58 and the cache 24 are interconnected to the native bus 46 and an internal device port access bus 64 for accessing the host port 18 and SCSI device ports 28 via a bus exchanger 66. The bus exchanger is a crossbar which provides fast access by all entities to all parts of the controller. The host port 18 may be implemented as a fast-wide differential SCSI-2 interface, using commercially available SCSI I/O processors and SCSI differential drivers. A SCSI I/O processor, having internal storage available to store binary-encoded ID values for the SCSI IDs, thus allows the controller to assume more than one host side SCSI bus address or ID. The SCSI device ports 28 can also be implemented by coupling commercially available SCSI port processor chips, which execute scripts read from the shared memory 58 under the control of the policy processor 40, with SCSI transceivers. Also residing on bus 64 is a RAID assist component 70, which accelerates RAID operations such as parity RAID parity calculations and host data compare requests.

Detailed Description Text (16):
FIG. 5 depicts the intercontroller communications interface 80 utilized by the storage controllers 14, 14' (shown here, along with their respective caches 24, 24', in dashed lines) in dual-active, redundant configuration during failover and failback operations. The interface between the controllers includes the communication link 22, implemented by means of a serial communication universal asynchronous receiver/transmitter (UART) on each controller. This communication link is used by each controller to send "keep alive" messages to the other controller. It is also used by the controllers to inform each other about the configuration information. The interface further includes CDAL buses 82, which are the data and addresses buses, as well as control signals, for accessing the caches. Not shown but also included are other signals which relate to the current state of the configuration and to specific control lines that determine the operation of the dual-active, redundant pair. These state and control signals are used by each controller to sense the presence or absence of the other controller in a dual-active redundant configuration and the status (e.g., presence or absence) of the other controller's cache, in addition to its own.

Previous Doc Next Doc Go to Doc#

L3: Entry 80 of 92 File: USPT Aug 4, 1998

US-PAT-NO: 5790775
DOCUMENT-IDENTIFIER: US 5790775 A

TITLE: Host transparent storage controller failover/failback of SCSI targets and associated units

DATE-ISSUED: August 4, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Marks; Randal S.	Colorado Springs	CO		
Roberson; Randy L.	Colorado Springs	CO		
Shen; Diana	Colorado Springs	CO		
Sicola; Stephen J.	Monument	CO		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Digital Equipment Corporation	Maynard	MA			02

APPL-NO: 08/ 546804 [PALM]
DATE FILED: October 23, 1995

INT-CL: [06] G06 F 11/20

US-CL-ISSUED: 395/182.07; 395/182.05
US-CL-CURRENT: 714/9; 714/7

FIELD-OF-SEARCH: 395/182.07, 395/182.05, 395/470, 395/183.04, 395/182.04, 733/135

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4589063	May 1986	Shah et al.	364/200
☐	5155845	October 1992	Beal et al.	395/575
☐	5274645	December 1993	Idleman et al.	371/10.1
☐	5276864	January 1994	Hrnandez et al.	395/182.08
☐	5319754	June 1994	Meinecke et al.	395/324
☐	5390187	February 1995	Stallmo	395/182.05
☐	5412661	May 1995	Hao et al.	395/182.04
☐	5471636	November 1995	Hauck	395/800

☐	5546535	August 1996	Stallmo et al.	395/182.07
☐	5561783	October 1996	Vanka et al.	395/468
☐	5581740	December 1996	Jones	395/500
☐	5586302	December 1996	Keener et al.	395/481
☐	5588111	December 1996	Cutts, Jr. et al.	395/182.07

ART-UNIT: 275

PRIMARY-EXAMINER: Beausoliel, Jr.; Robert W.

ASSISTANT-EXAMINER: Baderman; Scott T.

ATTY-AGENT-FIRM: Peterson; Cathy L. Hudgens; Ronald C.

ABSTRACT:

Provided herein is a method and apparatus for host transparent storage controller failover and failback. A controller is capable of assuming the identity of a failed controller while continuing to respond to its own SCSI ID or IDs in such a way that all SCSI IDs and associated units (LUNS) of the failed controller are effectively taken over by the surviving controller. This "failover" behavior is transparent to any attached host computers and is treated by such attached hosts as a powerfail condition. The symmetric operation of returning the targets (IDs) and units (LUNs) to the previously failing controller ("failback") is likewise transparent.

4 Claims, 9 Drawing figures

Previous Doc Next Doc Go to Doc#

L3: Entry 80 of 92 File: USPT Aug 4, 1998

US-PAT-NO: 5790775
DOCUMENT-IDENTIFIER: US 5790775 A

TITLE: Host transparent storage controller failover/failback of SCSI targets and associated units

DATE-ISSUED: August 4, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Marks; Randal S.	Colorado Springs	CO		
Roberson; Randy L.	Colorado Springs	CO		
Shen; Diana	Colorado Springs	CO		
Sicola; Stephen J.	Monument	CO		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Digital Equipment Corporation	Maynard	MA			02

APPL-NO: 08/ 546804 [PALM]
DATE FILED: October 23, 1995

INT-CL: [06] G06 F 11/20

US-CL-ISSUED: 395/182.07; 395/182.05
US-CL-CURRENT: 714/9; 714/7

FIELD-OF-SEARCH: 395/182.07, 395/182.05, 395/470, 395/183.04, 395/182.04, 733/135

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

Search Selected Search ALL Clear

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
☐	4589063	May 1986	Shah et al.	364/200
☐	5155845	October 1992	Beal et al.	395/575
☐	5274645	December 1993	Idleman et al.	371/10.1
☐	5276864	January 1994	Hrnandez et al.	395/182.08
☐	5319754	June 1994	Meinecke et al.	395/324
☐	5390187	February 1995	Stallmo	395/182.05
☐	5412661	May 1995	Hao et al.	395/182.04
☐	5471636	November 1995	Hauck	395/800

☐	5546535	August 1996	Stallmo et al.	395/182.07
☐	5561783	October 1996	Vanka et al.	395/468
☐	5581740	December 1996	Jones	395/500
☐	5586302	December 1996	Keener et al.	395/481
☐	5588111	December 1996	Cutts, Jr. et al.	395/182.07

ART-UNIT: 275

PRIMARY-EXAMINER: Beausoliel, Jr.; Robert W.

ASSISTANT-EXAMINER: Baderman; Scott T.

ATTY-AGENT-FIRM: Peterson; Cathy L. Hudgens; Ronald C.

ABSTRACT:

Provided herein is a method and apparatus for host transparent storage controller failover and failback. A controller is capable of assuming the identity of a failed controller while continuing to respond to its own SCSI ID or IDs in such a way that all SCSI IDs and associated units (LUNS) of the failed controller are effectively taken over by the surviving controller. This "failover" behavior is transparent to any attached host computers and is treated by such attached hosts as a powerfail condition. The symmetric operation of returning the targets (IDs) and units (LUNs) to the previously failing controller ("failback") is likewise transparent.

4 Claims, 9 Drawing figures

Previous Doc Next Doc Go to Doc#

L6: Entry 4 of 7 File: PGPB May 31, 2001

DOCUMENT-IDENTIFIER: US 20010002480 A1
TITLE: METHOD AND APPARATUS FOR PROVIDING CENTRALIZED INTELLIGENT CACHE BETWEEN MULTIPLE DATA CONTROLLING ELEMENTS

Abstract Paragraph:
Apparatus and methods which allow multiple storage controllers sharing access to common data storage devices in a data storage subsystem to access a centralized intelligent cache. The intelligent central cache provides substantial processing for storage management functions. In particular, the central cache of the present invention performs RAID management functions on behalf of the plurality of storage controllers including, for example, redundancy information (parity) generation and checking as well as AID geometry (striping) management. The plurality of storage controllers (also referred to herein as RAID controllers) transmit cache requests to the central cache controllers. The central cache controllers performs all operations related to storing supplied data in cache memory as well as posting such cached data to the storage array as required. The storage controllers are significantly simplified because the present invention obviates the need for duplicative local cache memory on each of the plurality of storage controllers. The storage subsystem of the present invention obviates the need for inter-controller communication for purposes of synchronizing local cache contents of the storage controllers. The storage subsystem of the present invention offers improved scalability in that the storage controllers are simplified as compared to those of prior designs. Addition of controllers to enhance subsystem performance is less costly than prior designs. The central cache controller may include a mirrored cache controller to enhance redundancy of the central cache controller. Communication between the cache controller and its mirror are performed over a dedicated communication link.

Detail Description Paragraph:
[0062] This central cache architecture improves overall subsystem performance by obviating the need for cache coordination message traffic over bus 208 thereby reducing overhead processing within the controller 204 and eliminating cache coordination message traffic over bus 208. Controllers 204 are therefore simpler than prior controllers exemplified as discussed above. The simpler controllers are substantially void of any local cache memory and parity assist circuits. The primary function served by the simpler controller is to provide an interface to attached host systems consistent with the storage management structure (e.g., RAID) of the subsystem. This simpler design permits easier scaling of the subsystem's performance by reducing the costs (complexity) associated with adding additional controllers. In like manner, additional intelligent central cache devices may be added either to increase the cache size and/or to provide mirrored redundancy of the central cache contents. As noted below with respect to FIG. 4, when adding cache devices to the central cache, it is preferred that the plurality of central cache devices communicate among themselves over a dedicated communication medium.

L6: Entry 4 of 7 File: PGPB May 31, 2001

PGPUB-DOCUMENT-NUMBER: 20010002480
PGPUB-FILING-TYPE: new-utility
DOCUMENT-IDENTIFIER: US 20010002480 A1

TITLE: METHOD AND APPARATUS FOR PROVIDING CENTRALIZED INTELLIGENT CACHE BETWEEN MULTIPLE DATA CONTROLLING ELEMENTS

PUBLICATION-DATE: May 31, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-47
DEKONING, RODNEY A.	WICHITA	KS	US	
WEBER, BRET S.	WICHITA	KS	US	

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	COUNTRY	TYPE CODE
LSI LOGIC CORPORATION				02

APPL-NO: 08/ 941770 [PALM]
DATE FILED: September 30, 1997

CONTINUED PROSECUTION APPLICATION: CPA

INT-CL: [07] G06 F 12/08

US-CL-PUBLISHED: 711/130; 711/113, 711/114, 709/328, 711/202, 711/135, 711/144
US-CL-CURRENT: 711/130; 711/113, 711/114, 711/135, 711/144, 711/202, 719/328

REPRESENTATIVE-FIGURES: 4

ABSTRACT:

Apparatus and methods which allow multiple storage controllers sharing access to common data storage devices in a data storage subsystem to access a centralized intelligent cache. The intelligent central cache provides substantial processing for storage management functions. In particular, the central cache of the present invention performs RAID management functions on behalf of the plurality of storage controllers including, for example, redundancy information (parity) generation and checking as well as AID geometry (striping) management. The plurality of storage controllers (also referred to herein as RAID controllers) transmit cache requests to the central cache controllers. The central cache controllers performs all operations related to storing supplied data in cache memory as well as posting such cached data to the storage array as required. The storage controllers are significantly simplified because the present invention obviates the need for duplicative local cache memory on each of the plurality of storage controllers. The storage subsystem of the present invention obviates the need for inter-controller communication for purposes of synchronizing local cache contents of the storage controllers. The storage subsystem of the present invention offers improved scalability in that the storage controllers are simplified as compared to those of prior designs. Addition of controllers to enhance subsystem performance is less costly than prior designs. The central cache controller may include a mirrored cache controller to enhance redundancy of the central cache controller. Communication between the cache controller and its mirror are performed over

a dedicated communication link.

RELATED PATENTS

[0001] This patent is related to commonly assigned, U.S. patent application Ser. No. 08/772,614 entitled METHODS AND APPARATUS FOR COORDINATING SHARED MULTIPLE RAID CONTROLLER ACCESS TO COMMON STORAGE DEVICES filed Dec. 23, 1996 which is hereby incorporated by reference.

Previous Doc Next Doc Go to Doc#